

Rec'd 11 APR 2005

Multicast Data Transfer

Field of the Invention

The invention relates to the simultaneous delivery of information to multiple hosts. In particular, the invention relates to an apparatus and method for transferring text, audio, video or other data to multiple hosts in a communication system.

Background

10 Methods for transmitting data from a source to multiple users over a network fall into one of three categories. Unicast is a point-to-point delivery mechanism, where a source transmits a separate stream of data for each user. The volume of data transmitted increases linearly with the number of users, leading to the consumption of an unacceptably high level of resources where the number of recipients is large.

15 The remaining two categories define point-to-multipoint or multipoint-to-multipoint delivery mechanisms. Broadcasting refers to the transmission of data to all hosts on a network. While this may work well in small networks, it requires the replication of a large number of data packets when used in a wide area network with many users. Furthermore, this method can be wasteful where the data is not of interest to a significant proportion of the recipients.

20 The third category is multicasting, which is the simultaneous transmission of data to a select group of users. Multicasting has particular application in fields requiring streaming of content to multiple users in real time, such as news feeds, online gaming, Digital Video Broadcasting (DVB) and videoconferencing. Overviews of this technique are given in "Multicast Networking and Application" by C. Kenneth Miller, Addison-Wesley 1988 [ISBN 0-201-30979-3] and in "Deploying IP MULTICAST in the Enterprise" by T. Maufer, Prentice Hall PTR, 1998 [ISBN 0-13-897687-2]. The use of multicasting in a cellular radio access network is discussed in "Multimedia Broadcast/ Multicast Service", 3rd Generation Partnership Project Technical Specification 3GPP™ TS 22.146 v1.0.0 2001.

Figure 1 is a block diagram of a prior communication system 1, for multicasting information from a content provider, i.e., host 2, to a plurality of receiver hosts 3 via a first network 4, such as the Internet, and a radio access network 5. For simplicity, a data path 6 through the Internet 4 is depicted as extending directly 5 from the content provider 2 to a router 7 without intervening stages, such as other transmitters, routers or servers, which may be located between these nodes. Data paths leading from the content provider 2 to the receiver hosts 3a-3c diverge at the router 7. Situations where this divergence occurs include those where the receiver hosts 3 do not belong to the same sub-network, where the receiver hosts 3 are 10 served by different transmitters 8 or where the radio access network 5 is not multicast enabled.

Instead of sending a separate stream of data packets A-D for each individual receiver host 3a, 3b, the content provider 2 provides a single data stream 6, which is 15 transmitted over the Internet 4 to the point at which the paths to the individual receiver hosts 3a, 3b diverge, which in this example comprises the router 7. At this point, the data stream of packets is replicated for onward transmission to the receiver hosts 3a, 3b individually, the router 7 being multicast-enabled in order to carry out the replication of the data. It will be understood that other routers (not 20 shown) in the network are also multicast enabled so that the data can be multicast to other receiver hosts individually. Thus, the router 7 copies incoming data packets A-E, producing duplicate data streams 9a, 9b for transmission to the receiver hosts 3a, 3b. As the content provider 2 transmits only a single data stream 6, the use of its system resources and the network load are reduced when compared 25 with a unicast system. Furthermore, as the transmission of replicated data is kept to a minimum, multicasting may be used in wide area networks where broadcasting is unfeasible.

There are two types of multicasting services. In the first, a content provider has a 30 reserved bandwidth for delivering a predetermined service to a group of end-users. The second type allows end-users to select the service provided to them. In both cases, a user can receive the multicast stream of data by submitting a request to join the group by sending a message to the content provider in a format specified by the

Internet Engineering Task Force (IETF). The user can leave the multicast group by submitting a corresponding request.

Multicasting is a suitable mechanism for providing end-users with a continuous service but has limited applicability for "one-shot" transmissions. For example, a group of users may require the delivery of a file, for example, a multimedia clip, a web page, an mp3 file, a document or a software module for an application or a game. Referring again to Figure 1, the users of receiver hosts 3a, 3b may both request the file, which comprises data packets A-E, from the content provider 2. As the data stream 6 is received by the router 7, the data packets are copied and forwarded to a radio transmitter 8, which transmits the encapsulated data packets to receiver hosts 3a, 3b as data streams 9a, 9b over a wireless network.

However, there are currently no provisions for allowing a host to join an ongoing file delivery transmission, so a third host 3c, requesting a file after the start of the transmission, would be excluded from the multicast file delivery transmission. Even if the third host 3c were to join the file delivery transmission, data packets A-C would have been missed and it is probable that the user of host 3c would require the entire file. This exclusion increases the likelihood that a transmission is received by a low number of recipients, or even a single user. Requests submitted after the delivery of the data packets A-E has commenced are accommodated by repeating the file delivery transmission.

Prior systems have addressed this problem by defining a schedule of forthcoming transmissions. A user can subscribe to a delivery list for a particular file, or set of files, and receive the relevant data at the next scheduled opportunity. However, this delivery mechanism does not respond directly to the user's requirements and is inflexible as, again, once the file delivery transmission has begun, any other further users requiring the same file are excluded and must subscribe to the next scheduled transmission.

In addition, a host 3b may drop one or more packets or, having received the first data packets A-C, become unable to receive the remaining packets D, E in a file

delivery transmission. The host 3b cannot make use of the data packets it has received and, in order to receive the entire file, data packets A-E, must wait for file delivery transmission, which may necessitate joining another multicast group.

5 The problems associated with "one-shot" transmissions are exacerbated in systems for multicasting data to mobile devices. In particular, a mobile device receiving the data packets over a cellular communications network may be located in a first cell when joining the group and move into a different cell during the file delivery transmission. The mobile device cannot receive the remaining data packets in the 10 file delivery transmission as the mobile device is no longer in the first cell.

The repetition of file delivery transmissions in order to accommodate further hosts and missed or dropped packets reduces the efficiency gain associated with multicasting and increases the load on the air interface bandwidth. As use of the air 15 interface for the provision of Internet type services rises, the efficient use of air interface bandwidth will become an increasingly important consideration.

Summary of the Invention

20 An object of the invention is to provide multicast file delivery in a manner that is more responsive to users' needs, with particular application to scalable multicasts such as unidirectional IP multicasts. The invention has particular advantages in systems for transferring data to mobile devices comprising a cellular 25 telecommunications network where a user may move between cells during the course of data transmission. Another object of the invention is to provide efficient use of the air interface bandwidth.

According to a first aspect of the invention, a multicast-enabled network element 30 comprises a first logical interface for receiving a file from a content provider, a second logical interface for forwarding said file to one or more hosts as a sequence of data packets in a file delivery transmission and a processor for defining a group of one or more hosts, wherein further hosts may be added to the group in response

to the reception of a request from said further host, the group being limited to further hosts situated within a single locational area, wherein said server is configured to transmit the file to the group so that said further hosts joining the group during transmission of the data packet sequence receive the remaining data 5 packets in said sequence.

According to a second aspect of the invention, a method of file delivery over a network comprises the steps of receiving at a network element a request for the file from a first host, retrieving the file from a content provider, defining a group 10 comprising the first host, forwarding the file to the group as a sequence of data packets in a file delivery transmission and adding to the group any further hosts submitting requests for the file during said file delivery transmission so that said further hosts receive the remaining data packets in said file delivery transmission, wherein the group is limited to further receiver hosts with a same locational area as 15 the first host.

This aspect further provides a computer program comprising program instructions for causing a network element to perform the method.

20 According to a third aspect of the invention, a method of file retrieval by a host over a network comprises the steps of sending a request to join a group to a network element, receiving a start packet transmitted by the network element which configures the connection between the network element and host, receiving a sequence of data packets transmitted by the network element in a first file delivery 25 transmission, arranging the sequence of data packets in their appropriate order and receiving a second file delivery transmission comprising the sequence of data packets, wherein the host retrieves any data packets that were dropped or missed in the first file delivery transmission by retrieving the corresponding data packets in the second file delivery transmission.

30

The invention allows a user to join a group, receive the remaining data packets in an ongoing file transmission and receive any missed data packets in the subsequent file transmission. Therefore, the time period between the submission of a file request

from a host and the completion of the file delivery is approximately the same for all hosts, although there will be small delays caused by the server sending additional packets for configuring connections. This compares favourably with the time period experienced in prior queuing systems, where the a host would have to wait 5 until a previous file delivery transmission has been completed, leading to improved data transmission rates and a quality of service that can be guaranteed.

Furthermore, the ability to join ongoing transmissions and receive the remaining data packets in a sequence assists in efficient delivery of a file where data packets 10 have been missed or dropped, or where the file delivery path to a receiver host changes during the transmission.

Brief Description of the Drawings

15 An embodiment of the invention will now be described with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a prior art communication system;
Figure 2 is a block diagram of a communication system in accordance with an 20 embodiment of the present invention at a first stage in a file delivery process;
Figure 3 is a flowchart of a multicasting process according to an embodiment of the present invention, as enacted by a server;
Figure 4 is a block diagram of part of the communication system shown in Figure 2 at a second stage in the file delivery process;
25 Figure 5 is a block diagram of part of the communication system shown in Figure 2 at a third stage in the file delivery process;
Figure 6 is a block diagram of a server in accordance with an embodiment of the present invention; and
Figure 7 depicts a mobile handset terminal for use according to an embodiment of 30 the invention.

Detailed Description

Figure 2 depicts an embodiment of a communication system 10 for multicast file delivery comprising a multimedia content provider 2, receiver hosts 3a-3c, a server 5, which includes means for performing the functions of a router, a radio transmitter 8 and a cache 11. The data paths leading from the content provider 2 to the receiver hosts 3a-3c have a common portion 6 extending between the content provider 2 and the server 7 over the Internet 4 and diverge at the server 7, which is the last network element in the data paths to the individual receiver hosts 3a-3c 10 before an air interface. The data path 6 may include further elements, such as other transmitters, routers or servers, which may be located between the content provider 2 and server 7.

In this particular example, receiver hosts 3a-3c are mobile devices belonging to a 15 cellular wireless DVB-T wide area network (WAN) 5. In Figure 2, the three receiver hosts 3a-3c are situated in an area covered by a single DVB-T cell C1, which is defined by a cell boundary 12 and associated with the transmitter 8. The area covered by the DVB-T cell C1 is used to define the location of the receiver hosts 3, so that, in this example, the receiver hosts 3a-3c have a common location. The 20 server 7 forwards the multicast data to the radio transmitter 8, which sends the data to the receiver hosts over the DVB-T network 5 in accordance with a suitable communication protocol. Suitable protocols include, but are not limited to, the following protocols: Reliable Multicast Transport Protocol (RMTP), Reliable Multicast File Transfer Protocol (RMFTP), Asynchronous Layered Coding (ALC), 25 NACK-Oriented Reliable Multicast (NORM), Pragmatic General Multicast (PGM), Tree Acknowledgement based protocol (TRACK), User Datagram Protocol (UDP) and Unidirectional Hypertext Transfer Protocol (UHTTP). The Session Announcement Protocol (SAP) may also be used for sending service information to the receiver hosts 3a-3c.

30

In this example, where the DVB-T network 5 is used for unidirectional communication only, communications between the receiver host 3a and the server 7, including file requests (REQ), are made via an alternative communication system,

such as a cellular telecommunications network 13. In the figures, the transmission of file data over the Internet 4 or DVB-T network 5 is indicated using solid lines, while communication over the telecommunications network 13 is shown with dashed lines. It is possible that, while belonging to a single DVB-T cell C1, the 5 receiver hosts 3a-3c may belong to different cells in the telecommunications network 13, or even to separate telecommunications networks, however, for simplicity, only a single telecommunications network 13 and associated transmitter 14 are depicted in the figures.

10 An embodiment of the file delivery procedure followed by the server 7 is described with reference to Figure 3. Beginning at step s0, the server 7 receives a request for a file, such as a multimedia clip, or a popular song in mp3 format, from a receiver host 3a (step s1) via the telecommunications network 13. The request from receiver host 3a includes information about its location, in terms of the DVB-T cell C1 in 15 which the receiver host 3a is located.

The server 7 requests the file from the content provider 2 (step s2), which responds by transmitting the file to the server 7, together with an associated delivery message. The delivery message includes information about the time the file was created and 20 an indication of when it will next be updated. The file and delivery message are stored in the cache 11.

The server 7 defines a group and sets up a corresponding file delivery list with entries that indicate receiver hosts 3 which have requested the file and are situated 25 within that locational area, i.e. within an area covered by the same DVB-T cell C1 (step s3). At this stage, the list contains only one entry, relating to receiver host 3a, although if a request for the same file is received from another receiver host (step s4) during retrieval of the file, the server 7 would check whether the new receiver host is located within the same DVB-T cell C1 as the first receiver host 3a (step s5). 30 If so, the new receiver host is added to the group at this stage by adding a corresponding entry to the file delivery list (step s6). If the new receiver host is at a second location, e.g. a second DVB-T cell C2, the server 7 may set up a separate file

delivery list for that location or, where a file delivery list for the second location C2 already exists, add the new receiver host to the relevant list.

5 The server 7 sends a “start” packet ST to the first receiver host 3a via the radio transmitter 8 (step s7). The start packet ST is also sent to any other receiver host currently in the group. The start packet ST is associated with an application that carries out the “handshake” procedure for establishing a connection between the server 7 and receiver host 3a.

10 Some communication protocols, e.g. NORM or PGM, require a receiver host 3a to send a negative acknowledgement (NACK) to the server 7 where a data packet has not been received or where a corrupted data packet has been received. A negative acknowledgement is a message requesting the retransmission of one or more missed data packets. Where the invention is applied in a system using such a protocol, the 15 connection between the server 7 and receiver host 3a is configured so that a period of time is allowed by the receiver host 3a before a negative acknowledgement message or status report is sent to the server 7. This period of time is longer than the time taken to receive the entire file. The reasons for this will be described later.

20 The file is then retrieved from the cache 11 and arranged into data packets A-G according to the appropriate communication protocol. The header information included in the data packets A-G includes information regarding its position in the order of the data packets within the file. The server 7 begins to forward copies of the data packets in sequence to the radio transmitter 8, where they are encapsulated 25 in an IP datagram and sent to the group (step s8). In this example, the group contains only the first receiver host 3a at this stage.

Unlike the prior system, the present multicast system allows further receiver hosts 3 to join the ongoing file delivery transmission. For example, if a file request from a 30 second receiver host 3b is received (step s9) while the file delivery transmission to receiver host 3a is in progress, e.g., after data packets A, B have been sent to the first receiver host 3a, the server 7 responds by checking whether the second receiver host 3b is at the same location as the first receiver host 3a, i.e. whether receiver host

3b is located in the area covered by DVB-T cell C1 (step s10). If so, the second receiver host 3b is added to the group (step s11) and a start packet ST is sent to the second receiver host 3b in order to establish a connection (step s12). The server 7 makes an entry in the file delivery list logging the next data packet A-G to be transmitted (step s13). This information is used to determine the point in the data packet sequence at which the second receiver host 3b joined the group. Subsequent data packets C-G are copied to the receiver hosts 3a, 3b in the group (step s8). As the data packet headers indicate their position in the data packet sequence, the second receiver host 3b, having missed the initial packets A, B can arrange the 10 incoming data packets C-G in the appropriate order in its IP stack.

Similarly, if a request is received from a third receiver host 3c within the same DVB-T cell C1 after data packets A-D have been transmitted, the server 7 will send a start packet ST to the third receiver host 3c and add the third receiver host 3c to 15 the group by including a corresponding entry to the file delivery list. Thus, the third receiver host 3c will receive data packets E-G during the ongoing file transmission.

When the last data packet in the file has been sent (step s14), the server 7 sends an 20 "end" packet END to the receiver hosts 3a, 3b in the group (step s15), which indicates that the end of the data packet sequence has been reached. The server 7 then checks whether the delivery of the file to each of the hosts 3a-3c has been completed (step s16). For example, the log associated with the file delivery list may be checked, so that the various points in the data packet sequence where the 25 receiver hosts 3a-3c joined the file delivery transmission can be used to determine whether any of the hosts 3a-3c require further data packets (step s16).

Alternatively, the connection between the server 7 and each receiver host 3a-3c may be configured so that it is terminated by either the server 7 or the receiver host 3a-3c and/ or the server 7 may remove the entry corresponding to that receiver host 3 30 from the file delivery list when the complete sequence of data packets A-G has been received. However, it may not be necessary to terminate the connection between the server 7 and receiver host 3a when the file delivery to that receiver host 3a has

been completed, as it is common practice to configure the receiver hosts 3 to drop any duplicate packets it receives.

The server 7 then allows a period of time for the receipt of requests for missing 5 packets from receiver hosts 3 that have not received the full sequence of data packets A-G. For example, in the situation shown in Figure 4, the first receiver host 3a has received the complete sequence of data packets A-G, while the second and third receiver hosts 3b, 3c have received only part of the data packet sequence C-G and F-G respectively, along with the start and end packets, ST, END. If one 10 or more receiver hosts 3b, 3c have not received the complete file, the server 7 continues to transmit the data packets A-G (step s8), beginning at the start of the data packet sequence. As described above, the receiver hosts 3b, 3c use the information in the data packet header regarding the position of the incoming data packet in the sequence to assemble the data packets in the correct order.

15

Furthermore, a receiver host 3 that did not receive a packet for any other reason, e.g., such as a problem with its connection to the server 7, can use the data packets sent in the retransmission of the file to complete the data packet sequence in its IP stack. For example, if receiver host 3b has received packet C, E-G but had dropped 20 packet D, the receiver host 3b can retrieve both the missing packets A, B and the dropped packet D from a subsequent file retransmission before the expiry of the time period allowed by the receiver host 3b before sending a negative acknowledgement. The use of data packets from subsequent file retransmissions reduces the number of negative acknowledgements and resulting retransmissions of 25 data packets to individual receiver hosts 3, thereby lessening network traffic.

However, the negative acknowledgement procedure would be used where a receiver host 3 has lost a packet in the final file transmission in the file delivery process due to transmission failure or corruption. In this particular case the receiver host 3a would send a negative acknowledgement message to the server 7 via the cellular 30 telecommunications network 13 requesting retransmission of the missing packets on the expiry of the time period.

Any negative acknowledgements that are sent by one of the receiver hosts, e.g. host 3c, are multicasted via the server 7 and radio transmitter 8 to the other receiver hosts 3a, 3b. This avoids the sending of multiple negative acknowledgements by the receiver hosts in relation to the same data packet.

5

A receiver host 3 may also miss data packets due to its movement out of the locational area associated with the group, i.e. out of the DVB-T cell C1 used to define the group location. For example, during the retransmission of data packet A shown in Figure 4, receiver host 3b may move over the cell boundary 12, in the 10 direction of arrow 15, to another DVB-T cell C2, as shown in Figure 5. While the server 7 will continue the file delivery transmission, transmitting data packets to receiver host 3c, receiver host 3b can no longer receive data stream 9b. After the expiry of the allowed time period, the receiver host 3b will send one or more negative acknowledgement messages relating to data packets A, B via 15 communications system 13, which include information regarding its new location. This will be received at a server 7' associated with the new DVB-T cell C2, which treats the NACK as a request for the file. The server 7' then follows the procedure of Figure 3, by requesting the file from the content provider 2, receiving the file via the Internet 4 over data path 6', storing it in a cache 11', setting up a group 20 associated with the location C2 of the receiver host 3b and transmitting the file via transmitter 8' as data stream 9b'. However, if there is an ongoing file delivery transmission for the same file at the new DVB-T cell C2, the server 7' will add the receiver host 3b to that group. The server 7 associated with the original DVB-T cell C1 could then delete the receiver host 3b from its group as, in the absence of any 25 negative acknowledgement messages, the file delivery to receiver host 3b would be considered complete, as described above in relation to step s16. The servers 7, 7', transmitters 8, 8' and caches 11, 11' may be physically provided as a single server 7, transmitter 8 and cache 11 configured to serve both cells C1, C2.

30 Once the file delivery has been completed, the file in cache 11 is deleted (step s17), completing the file delivery process (step s18). Alternatively, the deletion of the file in the cache 11 may be delayed for a predetermined period of time after the

completion of file delivery, in order to meet any further requests for the file without repeating the step of retrieving the file from the content provider 2 (step s2).

Under this mode of operation, after receiving a request for the file, the server 7 checks whether a copy of the file is already stored in its cache 11 and uses the

5 information in the delivery message to determine whether the file in cache 11 is the latest available version and whether it is necessary to retrieve an updated file from the content provider 2. This procedure may also be followed by the second server 7' mentioned above.

10 An example of a suitable server 7 is shown in detail in Figure 6. In addition to the cache 11, the server 7 comprises an input/output interface 16 for receiving data from the content provider 2 and input/output interfaces 17-19 for forwarding data to the receiver hosts 3a-3c. The file requests are received via the input/output interface 16, where file requests are submitted via a telecommunications network 13, 15 as in the figures, or via the Internet 4. Alternatively, where the wireless network 5 is used for bi-directional communication, the file requests may be received at interfaces 17-19. The interfaces 16-19 and a number of other components are connected by a data bus 20 as follows. A processor 21, with an associated clock 22, monitors and controls routing operations, including generating processes for 20 establishing and releasing connections between the server 7 and receiver hosts 3a-3c. A storage device 23 is provided for storing the server application software. An address processor 24 extracts the address information from incoming data packets and uses this information to determine how each data packet is to be processed, e.g., whether an incoming packet is a request for a file from a receiver host 3a-3c or 25 file data. Incoming file data is stored in the cache 11, although the server may also have access to an external buffer for storing large files. File requests are forwarded to a file request handler 25, which manages the copying of data packets A-G for transmission to the receiver hosts 3a-3c and the encryption of information relating to the correct order of the data packets A-G in the packet headers. The file request 30 handler 25 also maintains the file delivery list.

The receiver hosts 3 may be mobile telephone handset terminals, such as that shown in Figure 7. The mobile telephone handset terminal 26 comprises an antenna 27, a

transceiver 28, a central processing unit (CPU) 29 and a battery 30. The transceiver 28, which is, for example, a GSM, GPRS, 3G, or similar bi-directional transceiver, is connected to the antenna 27 and the CPU 29. Also connected to the CPU 29 are a receiver (DVB-T) 31, a keypad 32 and a display 33. The terminal 26 also includes 5 other conventional features of mobile telephone handsets, but these are omitted from Figure 7 for the sake of clarity.

The time interval between the submission of a file request by a receiver host and the receipt of the full sequence of data packets A-G should be the same for each of the 10 receiver hosts 3a-3c, allowing for small differences due to the transmission of further start packets ST to other receiver hosts 3. This predictability allows a service provider to guarantee the quality of service in terms of data transmission rates. This time interval compares favourably with a unicast system, or the prior multicast system where the available window for further receiver hosts 3b, 3c to join 15 an ongoing file delivery transmission is small. This provides potential efficiency gains as the likelihood that a data packet A in a file delivery transmission is sent to a single receiver host is decreased, which reduces the number of file delivery transmissions needed to accommodate a given number of receiver hosts 3.

20 The improvement in performance between the invention and a prior system to the receiver hosts during an example file delivery process is shown in Tables 1 and 2. The tables relate to the transmission of a 10 Mbit file, where the file delivery transmission to receiver host 3a commences at 0 s. The second and third receiver hosts 3b, 3c join the file delivery transmission at 10 s and 20 s respectively.

25 The information in Tables 1 and 2 presumes that no packets are dropped in the transmission, i.e., the data for the time of file delivery completion t_d is the minimum time for the reception of the complete sequence of data packets A-G by a receiver host 3. The values of t_d are also approximate as they do not include small delays 30 caused by sending the receiver hosts 3 start and end packets ST, END.

Table 1 illustrates an example where the system allocates a bandwidth of 100 kbit/ s per user. As the same data packets are transmitted simultaneously to the receiver

hosts 3a-3c, the available bandwidth is effectively shared between the receiver hosts 3a-3c. Therefore, as further receiver hosts 3a-3c are included in the file delivery transmission, the data transfer rate experienced by each of the receiver hosts 3a-3c increases. For example, the first receiver host 3a receives data at 100 kbit/ s over the time interval 0-10 s, i.e., until the second receiver host 3b joins the group. Data is then transferred to the receiver host 3a at 200 kbit/ s during the time interval 10-20 s. When receiver host 3c joins the group, the data transfer rate increases again to 300 kbit/ s.

				Prior System		Present system	
User	Time of Request (s)	File Size (Mbit)	Transfer Rate (kbit/s)	t_d (s)	PBW (kbit/s)	t_d (s)	PBW (kbit/s)
1	0	10	100	100	100	43	233
2	10	10	100	190	53	39	256
3	20	10	100	280	36	49	204

10

Table 1

In this example, it has been presumed that the transmission to each receiver host 3a-3c stops when the complete file has been received, i.e., that the data transfer rate 15 drops to 200 kbit/ s at 43 seconds, as the receiver host 3a has left the group. If the server 7 is configured to continue transmitting data to receiver host 3a until each host in the group has received the complete sequence of data packets A-G, the data transfer rate and perceived bandwidth will remain at 300 kbit/ s.

20 In the prior system, where further receiver hosts 3 cannot join an ongoing file delivery transmission, receiver hosts 3a, 3b and 3c would receive the file 100 s, 190 s and 280 s respectively, from the time that the file request from each receiver host 3 is submitted. The end users of the receiver hosts 3a-3c perceive bandwidths of 100 kbit/ s, 53 kbit/ s and 36 kbit/ s respectively. In the present system, the time 25 between submitting a request and completing the file delivery is 43 s, 39 s and 49 s respectively, so that the receiver hosts perceive bandwidths of 233 kbit/ s, 256 kbit/ s and 204 kbit/ s.

Table 2 compares the file delivery times and bandwidths perceived by end-users of receiver hosts 3 for the prior queue-based system and a system according to the invention in which the total allocated bandwidth is constant.

5

				Prior System		Present system	
User	Time of Request (s)	File Size (Mbit)	Transfer Rate (kbit/s)	t_d (s)	PBW (kbit/s)	t_d (s)	PBW (kbit/s)
1	0	10	100	100	100	100	100
2	10	10	100	190	53	100	100
3	20	10	100	280	36	100	100

Table 2

In this example, each of the receiver hosts 3a-3c receive the file in 100 s, without 10 allowing for dropped packets and the extra start and end packets ST, END, sent during the first file delivery transmission. Each receiver host 3a-3c perceives the total available bandwidth of 100 kbit/s.

The invention has been described by way of an example and can be used for 15 transferring data packets other than IP packets. While a point-to-multipoint transmission has been described, the invention may also be used for multipoint-to-multipoint data transfer. The receiver hosts 3a-3c may be fixed or mobile devices, e.g., mobile telephones or personal digital assistants (PDAs). It is not necessary for the receiver hosts 3a-3c to belong to the same local network or for the location of 20 the hosts to be defined in terms of cell coverage.

Furthermore, it is not necessary for the server 7 to transmit data to the receiver hosts 3a-3c using a wireless communication system. The embodiment described relates to a system comprising a WAN, however, the invention is equally applicable 25 to LAN networks or any other multicast enabled network and may be used for the transfer of files over networks other than the Internet or a DVB-T network, e.g., over DVB-S (satellite), DVB-C (cable), DVB-M, other DVB variants, Integrated

Services Digital Broadcasting (ISDB), ATSC Digital Television or Digital Audio Broadcasting (DAB) networks or in multicast-enabled General Packet Radio Service (GPRS), Enhanced Data GSM Environment (EDGE), Universal Mobile Telecommunications Systems (UMTS), Wireless Code Division Multiple Access (W-CDMA), CDMA2000 or Bluetooth systems.

The cellular telecommunications network 13 is used merely as an alternative communication system to the wireless system 5. A different system, such as the Internet 4, could be used to convey requests and negative acknowledgements from the receiver hosts 3 to the server 7. If the wireless system 5 is bi-directional, the invention could be implemented using a single system to handle file delivery to the receiver hosts 3 and file requests.

In the embodiment described above, the location associated with the groups was defined in terms of the area covered by a DVB-T cell C1. However, the location may be defined differently, as required, for example in terms of the cells of the telecommunications network 13, depending on the requirements of the user or a service provider.

Furthermore, although the described embodiment comprised a server 7 located at the air interface, the invention may be implemented at any network element, e.g., a server, node or router, along the path between a content provider 2 and receiver hosts 3a-3c, while the receiver hosts may be fixed or mobile devices, including mobile personal digital assistants (PDAs).